

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.



12/31/98
PTO
U.S.



UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 003813.P003

Total Pages 5

12/31/98
PTO
U.S.

First Named Inventor or Application Identifier Raghunath Iyer

Express Mail Label No. EL164802299US

ADDRESS TO: **Assistant Commissioner for Patents**
Box Patent Application
Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)
2. Specification (Total Pages 31)
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claims
 - Abstract of the Disclosure
3. Drawings(s) (35 USC 113) (Total Sheets 8)
4. Oath or Declaration (Total Pages 5 unsigned)
 - a. Newly Executed (Original or Copy)
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed) (**Note Box 5 below**)
 - i. **DELETIONS OF INVENTOR(S)** Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (useable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Microfiche Computer Program (Appendix)

7. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
a. Computer Readable Copy
b. Paper Copy (identical to computer copy)
c. Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. Assignment Papers (cover sheet & documents(s))
9. a. 37 CFR 3.73(b) Statement (where there is an assignee)
 b. Power of Attorney
10. English Translation Document (if applicable)
11. a. Information Disclosure Statement (IDS)/PTO-1449
 b. Copies of IDS Citations
12. Preliminary Amendment
13. Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
14. a. Small Entity Statement(s)
 b. Statement filed in prior application, Status still proper and desired
15. Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. Other: Express Certificate of Mailing

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:

Continuation Divisional Continuation-in-part (CIP)
of prior application No: _____

18. **Correspondence Address**

Customer Number or Bar Code Label _____
(Insert Customer No. or Attach Bar Code Label here)
or
 Correspondence Address Below

NAME Dag H. Johansen (Reg. No. 36,172)
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard
Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

Country U.S.A. TELEPHONE (408) 720-8598 FAX (408) 720-9397

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**Method and Apparatus for High-Speed
Network Rule Processing**

Inventors: Raghunath Iyer
Sundar Iyer
Moti Jiandani
Ramana Rao

Prepared by:
Blakely, Sokoloff, Taylor & Zafman
1279 Oakmead Parkway
Sunnyvale, California 94086
(408) 720-8598

"Express Mail" mailing label number EL164802288US
Date of Deposit 12/31/98

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Grace A. Cruz
(Typed or printed name of person mailing paper or fee)

Grace A. Cruz
(Signature of person mailing paper or fee)

Method and Apparatus for High-Speed Network Rule Processing

FIELD OF THE INVENTION

5

The present invention relates to the field of computer networking. In particular the present invention discloses a method and apparatus for quickly processing packets that are tested with a large number of rules.

BACKGROUND OF THE INVENTION

The Internet is a global interconnection of computer networks that share a set of well-defined data communication protocols. Specifically, most computer networks coupled to the global Internet communicate using the Transport Control Protocol (TCP) 15 and Internet Protocol (IP).

A very large portion of the computers communicate on the global Internet are coupled to a local area network (LAN) that is coupled to the global Internet with an Internet gateway. The Internet gateway handles all communication between computers 20 on the local area network and computers out on the global Internet. The Internet gateway may perform many different functions such as network address translation, network caching, routing, and packet filtering.

Packet filtering is the task of examining each packet to apply a set of 25 filtering rules. Each packet filtering rule specifies a particular packet filtering policy. For

example, all packets incoming from the Internet that are destined for vulnerable server ports may be discarded in order to protect the internal servers on the local area network.

The number of packet filtering rules that are needed depends on the

- 5 particular application. In simple packet filtering routers for small or home routers, the number of packet filtering rules is relatively small. However, an internet service provider (ISP) that provides classes of service for the internet service provider's customers, the internet service provider will need many thousands of packet filter rules to implement the class of service priority and other customer features.

10

There are several different current implementations of packet filtering rule processors. The simplest implementation of a rules processor is a linear searching rule processor. In such an implementation, the linear rule processor tests each received packet against each rule in the list of packet filtering rules. The time required to perform this 15 type of rule processing is directly proportional to the number of packet filtering rules. This type of linear rule processing is not feasible for any system with a large number of packet filtering rules.

To provide faster rule processing, improved methods of applying packet

- 20 filtering rules were introduced. One improved method is known as "rule splitting". A rule splitting system divides the rules into several different sets of rules. When a packet is received, one or more aspects of the packet are examined to determine which subset of rules should be applied. For example, a rule splitting type of rule processor may only examine the Source and/or Destination ports of each packet in order to determine which

set of rules to apply. Rule splitting type of rule processors are difficult to implement because rules have ranges associated with them.

Search trees provide another method of improving rule processing speed.

- 5 Search trees divide the rules into a preprocessed organized format that improves the rule processing to a speed that is a logarithmic function of the number of rules in each dimension. This type of searching works well, but still does not provide a solution that is viable for high-speed network applications that require many thousands of rules. It would therefore be desirable to have an improved network rule processor that can process
- 10 thousands of network rules.

03813.P003

SUMMARY OF THE INVENTION

The present invention introduces a high-speed rule processing method that may be used for packet filtering. The high-speed rule processor pre-processes a set of 5 packet filtering rules such that the rules may be searched in parallel by a set of independent search units.

In the rule pre-processing of the method of the present invention, a set of packet filtering rules are first divided the rules into N dimensions. The N dimensions are 10 orthogonal aspects of each packet that may be examined and tested in each rule. Each of the N dimensions are then divided into a set of dimension rule ranges wherein each rule range defines a non-overlapping contiguous range of values in a particular dimension and the rules that may apply to packets that fall within that rule range. Each rule range may be assigned an R-length bit vector that specifies the rules that may apply to packets that 15 fall within that rule range. If the rules are prioritized wherein only the highest priority rule will be applied then such bit vectors will be organized into an order bit vector wherein the highest priority rule is at the beginning of the rule bit vector and the lowest priority rule will be at the end of the rule bit vector. The rule preprocessing is completed by creating a search structure (such as a look-up table, Patricia tree structure, or binary 20 tree structure) for each of the N dimensions. Each search structure may be used by an independent search unit such that all N dimensions may be searched concurrently.

The packet processing method of the present invention activates the N independent search units to search the N pre-processor created search structures. In one 25 embodiment, the output of each of the N search structures will be an R-length bit vector.

In such an embodiment, the N output bit vectors are logically ANDed together to produce a final rule bit vector that is used to select the rule or rules to be applied. If the rules are prioritized, then only the first matching rule (the highest priority rule) will be applied.

5 Other objects, features, and advantages of present invention will be apparent from the company drawings and from the following detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the present invention will be apparent to one skilled in the art, in view of the following detailed description in which:

5

Figure 1 illustrates a typical small local area network (LAN) coupled to the Internet through a packet filtering device that filters the packets that travel between the global Internet and the local area network.

10

Figure 2 illustrates the format of a packet from the Transport Control Protocol layer.

Figure 3 illustrates the format of a packet from the Internet Protocol layer.

15

Figure 4 illustrates a two dimensional rule space that contains two different rules.

Figure 5 illustrates the two dimensional rule space of **Figure 4** wherein each dimension has been divided into different rule ranges.

20

Figure 6 illustrates the two dimensional rule space of **Figure 5** wherein each rule range has been assigned a bit vector that specifies the rules that may apply.

25

Figure 7 illustrates a tree structure that may be used to organize the X aspect ranges along the X dimension illustrated in **Figures 4 to 6**.

Figure 8 illustrates a flow diagram that summarizes the rule processing methods of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

A method and apparatus for performing high-speed network rule processing is disclosed. In the following description, for purposes of explanation, 5 specific nomenclature is set forth to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention. For example, the present invention has been described with reference to an embodiment within a gateway that coupled a local area network to the global Internet. However, the same rule processing 10 engine can easily be used in other applications such as general purpose routers, firewalls, and routers that give priority to packets having a higher class of service. Furthermore, the rule processing teachings of the present invention have been disclosed in a packet filtering application although the rule processing teachings may be used in many other contexts.

15

Network Packet Filtering

Network packet filtering is the task of examining each packet received from a network and then applying a set of packet filtering rules to the received packets. 20 There are a number of uses for network packet filtering. For example, packet filtering can be used to provide security for a local area network by filtering out packets from potential intruders.

Example Packet Filtering Applications

Figure 1 illustrates one possible network application of packet filtering.

Referring to **Figure 1**, a local area network (LAN) **140** is illustrated with a set of computer workstations **142, 144, 146, and 148** coupled thereto. The local area network 5 (LAN) **140** environment also includes a couple of servers **141** and **143**. If the servers **141** and **143** are only used for internal purposes then the internal servers **141** and **143** should only be accessible by internal clients such as computer workstations **142, 144, 146, and 148**.

10 All communication between computer systems on the internal local area network **140** and the global Internet **100** passes through an Internet gateway **130**. The Internet gateway **130** may comprise a suite of firewall applications on a computer system, a packet filtering router, or another type of network component that provides the desired features. Of particular interest to the present invention, most Internet gateway systems 15 provide packet filtering features.

As specified earlier, an important packet filtering application is network security. Specifically, when a network node on the global Internet **100** attempts to establish a connection to an internal server (such as server **141**) on the local area network 20 **140** then the Internet gateway **130** should carefully scrutinize the associate network packets since the entity attempting to establish the connection may be attempting an unauthorized access. To prevent such unauthorized access, the Internet gateway **130** processes packets with a set of security rules that screen out packets related to unauthorized actions. For example, if the servers **141** and **143** are only to be used for 25 internal purposes then the Internet gateway **130** should screen out all packets originating

from the global Internet **100** and destined for the internal servers **141** and **143**. Such screening will be performed by the packet filtering rule processor of the Internet gateway **130**.

5 There are many different permutations of network security rule policies that may be implemented. For example, if the servers **141** and **143** are only to be used for internal purposes for one set of services (such as an internal database) and global access for other services (such as email) then the packet filtering rules in the Internet gateway **130** should screen out packets originating from the global Internet **100** and destined for
10 the internal servers **141** and **143** only if the service requested is for internal use only. In a TCP/IP environment, the service is usually specified using a TCP port number.

Packet filtering can also be used by an Internet service provider to provide various different classes of service to a set of users. For example, packets from a high priority user will be processed and transmitted faster than packets from a low priority user. With such a packet filtering arrangement, the Internet service provider can charge different prices for different classes of service. There may be many different classes of service offered by the Internet service provider. For example, an Internet service provider may provide a large scale of different services from a free (advertising subsidized) email-
20 only service to a high-speed high-quality service that can be used for low-latency tasks such as video conferencing. Such an application would require a very large number of rules applied to each packet to determine how the packet should be handled (if at all). Current packet filtering implementations are not capable of processing packets with a very large number packet filtering rules with low latency.

25

Packet Filtering Aspects

Network packets can be filtered by examining a number of different packet aspects. In an Internet application, the most interesting aspects to examine are in the Transport Control Protocol (TCP) packet header (or a UDP packet header) and the

- 5 Internet Protocol (IP) packet header. In normal Internet usage, TCP packets are enclosed within IP packets.

Figure 2 illustrates the format of a Transport Control Protocol (TCP)

packet. The objective of the Transport Control Protocol layer is to ensure a reliable

- 10 connection between two computer nodes. Of particular interest for packet filtering are the source TCP port **210** and the destination TCP port **220**. In some applications, a packet filtering device may examine the flags **230** in the TCP packet header.

The TCP packet is typically enclosed within an Internet Protocol packet.

- 15 **Figure 3** illustrates an Internet Protocol (IP) packet header. The objective of the Internet Protocol layer is to ensure that an Internet packet is properly routed from a source node having a source Internet Protocol (IP) address **310** to a destination node having a destination Internet Protocol (IP) address **320**. Packet filtering devices examine the source Internet Protocol (IP) address **310** and the destination Internet Protocol (IP) address **320**. Another important aspect that may be examined is the Protocol **330** field in the IP packet header. The Protocol **330** field specifies the type of protocol being used to transport the packet. If a TCP packet is enclosed, then the Protocol will specify the TCP protocol. Finally, a packet filter may also examine the flags **340** in the Internet Protocol header.
- 20

25

The Rule Processing Task

The task of rule processing consists of examining a set of aspects for a particular packet and determining the first rule that applies to a packet having those aspects. Each packet filter rule can be expressed as a set of ranges for each packet aspect
5 that will be examined. If all of the packets aspects fit within the defined ranges, then the rule will be applied. When a particular aspect is irrelevant for a particular rule, the rule may specify the full possible range. When a particular aspect must match a certain value exactly, the range for that particular aspect begins and ends with that certain value. Since more than one rule polytope may encompass a particular packet aspect point, the rule
10 filter must select the rule having the higher priority.

Using a set of aspect ranges, packet filtering becomes a problem in computational geometry. Specifically, each packet filtering rule defines a polytope in K dimensions wherein K is the number of packet aspects examined. Thus, given a point
15 defined by the aspects of a packet, the packet filtering engine must find the set of polytopes that encompass the point. For example, **Figure 4** illustrates a pair of two dimensional polytopes that define two different rules in a two aspect (dimension) rule space.

20 In the example of **Figure 4**, each aspect is a positive or zero integer. The aspects may be addresses, ports, or other content from received packets. Since there are two aspects that are examined in the example of **Figure 4**, the rule space is a two aspect/dimension rule space wherein each rule defines a rectangle polytope. Thus, Rule A forms a first rectangle and Rule B forms a second rectangle. A third

aspect/dimensional would create a rectangular solid. Further dimensions are not easily visualized but easily implemented in practice.

High-Speed Packet Filter Rule Processing

5

To improve the performance of a firewall, router, switch, or other network component that performs packet filtering, the present invention introduces a high-speed rule processing engine. The high-speed rule processing engine pre-processes the rules to divide the rules into different packet aspect dimensions. Each dimension is then divided
10 into individual rule ranges wherein each range has a different set of rules that may apply when compared to the adjacent ranges. Each range is then assigned a bit vector that specifies the rules that may be activate when a packet falls within the range. The pre-processing is completed by creating a different data structure to be used for searching each different dimension range. Examples of possible data structures include look-up
15 tables and organized data trees. When a packet is received, all of the search data structures for each aspect dimension are examined in parallel to generate the assigned bit vector for that range. The bit vector results of the parallel rule processors are combined into a final rule or set of rules that need to be applied to the packet.

20 **Rule Dimension Ranges**

To implement a high-speed rule processor for packet filtering, each dimension is divided into distinct dimension rule ranges wherein each distinct dimension rule range adds or omits a possible rule from an adjacent range. For example, **Figure 5** illustrates the two dimensional rule space of **Figure 4** after both dimensions have been
25 divided into distinct dimension rule ranges.

Referring to the X aspect/dimension of **Figure 5**, in the range from 0 to 9, no rule applies. In the range from 10 to 19 along the X aspect/dimension, Rule A may apply. (Note that, the Y aspect/dimension must be examined to determine if Rule A really does apply.) In the range from 20 to 29 along the X aspect/dimension, Rules A and/or B may apply. In the range from 30 to 39 along the X aspect/dimension, Rule B may apply. Finally, in the range above 40 in the X aspect/dimension, no rule applies.

The Y aspect/dimension of **Figure 5** is also divided into distinct dimension rule ranges. Specifically, from 0 to 9 in the Y aspect/dimension no rule applies, from 10 to 19 Rule B may apply, from 20 to 24 rules A and/or B may apply, from 25 to 29 rule A may apply, and above 30 no rule applies.

Rule Dimension Range Bit Vectors

After each aspect/dimension has been divided into distinct dimension rule ranges, each dimension rule range may be assigned an R-length bit vector that specifies the rules that apply to that particular dimension rule range. The R value specifies the number of rules that are being applied. If the rules are ordered in priority, then the R-length rule bit vector should be ordered in the same manner. **Figure 6** illustrates the two dimensional rule space of **Figure 5** wherein each dimension rule range has been assigned a bit vector. Since there are only two rules, the bit vector is 2 bits in length. The first position in the bit vector designates if points within the range may require Rule A. The second position in the bit vector designates if points within the range may require Rule B.

Referring to **Figure 6**, the first range where no rule applies is assigned "00". The second range along the X aspect/dimension wherein Rule A may apply is assigned "10" since Rule A (the first bit in the bit vector) may apply. The third range along the X aspect/dimension wherein both Rule A and Rule B may apply is assigned 5 "11" since both Rule A (the first bit in the bit vector) and Rule A (the second bit in the bit vector) may apply. The remainder of the ranges along the X aspect/dimension are assigned bit vectors in the same manner. The ranges along the Y aspect/dimension are also filled in the same manner.

10 To use the assigned bit vectors, a packet filtering engine first classifies a received packet into the appropriate range within each dimension. In the example of **Figure 6**, each packet is classified into one range along the X aspect/dimension and another range along the Y aspect/dimension. For example, a packet with an X aspect of 23 and Y aspect of 17, written as P(23,17), would be placed into the third range (from the 15 left) along the X aspect/dimension and the second range (from the bottom) along the Y aspect/dimension. The assigned bit vectors from the different dimensions are then logically ANDed together. In the previous example of P(23,17), the "11" bit vector from the X aspect/dimension is logically ANDed with the "01" bit vector from the Y aspect/dimension to generate "01". The "01" resultant vector specifies that the second rule, Rule B, should be applied. For clarity, the ANDed values of all the different range 20 combinations have been placed onto the rule space of **Figure 6**.

If an incoming packet has an X aspect of 27 and Y aspect of 21, P(27,21), then the packet would be placed into the third range (from the left) along the X aspect/dimension and the third range (from the bottom) along the Y aspect/dimension. 25

These classifications would output the bit vectors "11" in the X range and "11" in the Y range that are logically ANDed to produce "11". The "11" output bit vector means that both Rule A and Rule B apply to the packet. However, if the rules have been ordered in priority then only the first rule in the bit vector, the first one ("1"), will be used such that

5 only Rule A will be applied.

Alternate Embodiment: rule range look-up

In an alternate embodiment, each range may simply be assigned a range value. For example, referring to the X aspect/dimension of **Figure 4 to 6**, the range from

10 0 to 9 could be assigned a range value of 1, the range from 10 to 19 could be assigned the range value of 2, the range from 20 to 29 could be assigned the range value of 3, the range from 30 to 39 could be assigned the range value of 4, and the range above 40 could be assigned the range value of 5. The Y aspect/dimension would be assigned range values in a similar manner. Specifically, in the Y aspect/dimension, from 0 to 9 is rule range 1, 10 to 19 is rule range 2, from 20 to 24 is rule range 3, from 25 to 29 is rule range 4, and above 30 is rule range 5. These range values could then be used to look up a final rule to apply in an N dimensional look-up table. In the case that only the highest priority rule for a given field is needed, that value can be pre-computed amongst all matching rules in a given row and only that rule can be stored for each row in the table. In an embodiment

15 without a pre-computed highest priority rule, the table for the example in **Figure 4 to 6** might appear as follows:

Table 1

Packet Aspect Range Value in X and Y dimension	Output rule vector
1, 1	00
1, 2	00
1, 3	00
1, 4	00
1, 5	00
2, 1	00
2, 2	00
2, 3	10
2, 4	10
2, 5	00
3, 1	00
3, 2	01
3, 3	11
3, 4	10
3, 5	10
4, 1	00
4, 2	01
4, 3	01
4, 5	00
5, 1	00
5, 2	00
.	00
.	
.	

Such a table based system might be useful for systems with small numbers of rules. However the memory that would be required is proportional to two times the number of Rules plus one raised to the power of the number of dimensions, expressed mathematically as $(2^{\text{Rules}} + 1)^{(\text{Number of dimensions})}$. Thus, the memory requirements grow exponentially. Thus, such a system is not feasible for systems with large numbers of rules.

Rule Dimension Range Classification

As previously set forth, each received packet must be classified into a range along each examined packet aspect/dimension. Since these classifications are data independent along the different dimensions, the classification along each dimension may
5 be performed in parallel by individual logic modules. In this manner, a very complex search may be performed by N logical searching units concurrently wherein N indicates the number of dimensions. With independent parallel searching, a single very fast search module is replaced with N search modules with lower clock speeds. The time complexity in such a parallel implementation is reduced to order of $\log(2R + 1)$ for a binary tree
10 search on each dimension. Then an additional time is required to calculate the bit-wise logical AND of these results. This method is faster than all previously known implementations.

The rule range classification along each dimension is still a time
15 consuming problem. Every new rule raises the amount of time needed for searching each dimension. Specifically, there will be up to $2R + 1$ rule ranges in each dimension. ($2R + 1$ represents the worst case scenario.) With a dimension that has m possible variations, the task is to classify the incoming packets having m possible variations into one of the
20 up to $2R + 1$ rule ranges. For example, an 8-bit aspect in a packet will have 256 possible variations.

A simple method of range classifying is to use a long list ranges that specify the required condition of each range. Such a method is simple to implement but requires far too much time to search.

One of the fastest methods of classifying the incoming packets is to create a look-up table structure that maps each possible dimension value into the appropriate bit vector value for the associated rule range. The following table illustrates one possible implementation of a look-up table for the X dimension of the rule space in **Figures 4 to 6**.

- 5 The table outputs the appropriate bit vector based upon the X aspect of an incoming packet.

Table 2

Packet Aspect Value in X dimension	Output bit vector
0	00
1	00
2	00
3	00
.	00
.	
.	
9	00
10	10
11	10
.	10
.	
.	
19	10
20	11
21	11
.	11
.	
.	
29	11
30	01
31	01
.	01
.	
.	
39	01
40	00
.	00
.	

The look-up table provides the fastest possible method of classifying a particular dimension into a range with an output bit vector. However, a look-up table is not always feasible. In a system that has a large number of rules and dimensions that

5 have large numbers of possible combinations, the memory required to store a look-up table becomes impractical. For example, a 32 bit IP address has 2^{32} possible permutations such that a very large 2^{32} entry look-up table would be required. Furthermore, in a system with thousands of rules, each bit vector would be thousands of bits long. Constructing such a large look-up table is not economically feasible.

10

Another method of creating a data structure for quickly generating bit vectors is to use a search tree structure such as a Patricia tree or a binary tree. A search tree structure greatly reduces the amount of memory required by eliminating redundant data from the search structure. However, a tree structure will require more processing

15 time to search. Specifically, such a tree search will require a time amount that is a logarithmic function of the two times the number of rules plus one. [Order($\log(2R + 1)$)]

Figure 7 illustrates one possible tree structure that can be used to quickly search the rule ranges of the X dimension illustrated in **Figures 4 to 6**.

20

Since the different aspects of a packet will vary in the number of possible combinations, each dimension can be implemented with a different search structure. For example, one possible packet filtering engine would examine a 32-bit source IP address, a 32-bit destination IP address, and an 8-bit protocol value. In such a three dimensional packet filtering system, the relatively small 8-bit protocol dimension could be classified

25 quickly with a two hundred and fifty-six (256) entry look-up table and the larger 32-bit

dimensions could be classified using Patricia tree structures or binary tree structures. All three different dimensions could be searched in parallel using three different searching units. The bit vector outputs from the three search units could then be logically ANDed to generate a final rule bit vector that specifies the rule (if rules are prioritized) or rules (if all matching rules are to be applied) to be used on the packet.

5

Summary

Figure 8 illustrates a flow chart that summarizes the teachings of the present invention. Referring step **810** of **Figure 8**, the first task is to divide the rules into

10 N aspects/dimensions. The N dimensions are orthogonal aspects of each packet that may be examined and tested in each rule. For example, in a packet filtering system that only examines and tests the 32-bit source IP address, a 32-bit IP destination address, the 16-bit source TCP port, and the 16-bit destination TCP port, the packet filtering engine would be divided into four dimensions.

15

Next, at step **820**, each of the N dimensions is divided into a set of dimension rule ranges. Each rule range defines a non-overlapping contiguous range of values in particular dimension and the rules that may apply to packets that fall within that rule range. At step **830**, each rule range is assigned an R-length bit vector that specifies 20 the rules that may apply to packets that fall within that rule range. If the rules are prioritized wherein only the highest priority rule will be applied then the bit vectors will be organized into an order bit vector wherein the highest priority rule is at the beginning of the rule bit vector and the lowest priority rule will be at the end of the rule bit vector.

Finally, at step **840**, the rule preprocessing is completed by creating a search structure (such as a look-up table, Patricia tree structure, or binary tree structure) for each of the N dimensions. The N search structures may be different for each dimension. Each search structure may be used by an independent search unit such that all N dimensions may be searched concurrently.

5 N dimensions may be searched concurrently.

At step **850**, the packet processing may begin. When a new packet is received at step **850**, the method proceeds to step **860** wherein the N dimensions are searched using the created search structures. The output of each of the N search structures will be an R-length bit vector. At step **870**, the N output bit vectors are logically ANDed together to produce a final rule bit vector. Finally at step **880**, the final rule bit vector is used to select the rule or rules to be applied. If the rules are prioritized, then only the first matching rule (the highest priority rule) will be applied. The method returns back to step **850** to process additional packets.

15

New rules may be added at step **850**. If a new rule is submitted, the method proceeds to step **890** wherein the rule is examined to determine if the new rule examines and tests a new aspect of the packet. If a new aspect is to be tested, then the method proceeds to step **810** wherein a new dimension will be added to the packet filtering engine. Alternatively, if at step **890** the method determines that the new rule will only test an existing packet aspect, then the method proceeds to steps **820** through **840** wherein the new rule is preprocessed to create a new set of rule ranges for each dimension, new longer rule bit vectors are generated, and new search structures are created.

25

The foregoing has described a method and apparatus for performing high-speed rule processing for packet filtering. It is contemplated that changes and modifications may be made by one of ordinary skill in the art, to the materials and arrangements of elements of the present invention without departing from the scope of the
5 invention.

CONFIDENTIAL

CLAIMS

We claim:

1 1. A method of performing network packet filtering, said method
2 comprising:
3 preprocessing a set of rules to generate a set of rule ranges along N dimensions;
4 searching said rule ranges along said N dimensions in parallel to generate N sets
5 of possible rules along said N dimensions;
6 logically combining said N sets of possible rules to generate a final set of possible
7 rules; and
8 applying said final set of possible rules.

1 2. The method as claimed in claim 1 further comprising:
2 generating a rule bit vector for each rule range along each of said N dimensions;

1 3. The method as claimed in claim 1 further comprising:
2 generating a search structure for each set of rule ranges along each of said N
3 dimensions;

1 4. The method as claimed in claim 3 wherein one of said search
2 structures comprises a look-up table.

1 5. The method as claimed in claim 3 wherein one of said search
2 structures comprises a tree search structure.

1 6. The method as claimed in claim 1 wherein applying said final set
2 of possible rules comprises selecting a highest priority rule in said final set of possible
3 rules.

1 7. The method as claimed in claim 1 wherein applying said final set
2 of possible rules comprises applying all rules in said final set of possible rules.

1 8. The method as claimed in claim 1 wherein each of said N sets of
2 possible rules comprise a rule bit vector that specifies a set of rules that may apply.

1 9. The method as claimed in claim 8 wherein said rule bit vectors are
2 logically ANDed together to produce a final bit vector of rules that apply.

1 10. The method as claimed in claim 9 wherein applying said final set
2 of possible rules comprises selecting a highest priority rule in said final set of possible
3 rules.

1 11. A method of pre-processing a set of rules for processing incoming
2 data units, said incoming data units having a set of N dimensions to examine, said method
3 comprising:

4 dividing each of said N dimensions into a contiguous set of rule ranges;
5 assigning each of said rule ranges a range identifier; and
6 creating a search structure for each of said N dimensions that organizes said rule
7 ranges along each dimension such that an incoming data unit may be quickly
8 classified into one of said rule ranges.

1 12. The method as claimed in claim 11 wherein said range identifier
2 comprises a rule bit vector that specifies a set of rules that may apply to incoming data
3 units that fall within the associated rule range.

1 13. The method as claimed in claim 12 wherein said rule bit vectors
2 are logically ANDed together by a rule processor to produce a final bit vector of rules that
3 apply.

1 14. The method as claimed in claim 13 wherein said rule processor
2 selects a highest priority rule in said final set of possible rules.

)

1 15. The method as claimed in claim 11 wherein said range identifier
2 comprises an index value.

1 16. The method as claimed in claim 15 wherein said index values are
2 used by a rule processor to index into a N dimensional look-up table for a final rule.

1 17. The method as claimed in claim 11 wherein one of said search
2 structures comprises a look-up table.

1 18. The method as claimed in claim 11 wherein one of said search
2 structures comprises a tree search structure.

1 19. A method of processing incoming data units, said incoming data
2 units having a set of N dimensions to examine, said method comprising:
3 searching N search structures for said N dimensions in parallel to classify
4 incoming data units into a matching rule range along each of said N
5 dimensions;
6 logically combining said N matching rule ranges to generate a final set of possible
7 rules; and
8 applying said final set of possible rules.

1 20. The method as claimed in claim 19 wherein each rule range
2 comprises a range identifier.

1 21. The method as claimed in claim 20 wherein said range identifier
2 comprises a rule bit vector that specifies a set of rules that may apply to incoming data
3 units that fall within the associated rule range.

1 22. The method as claimed in claim 21 wherein said step of logically
2 combining said N matching rule ranges comprises logically ANDing together said rule bit
3 vectors to produce a final set of possible rules.

1 23. The method as claimed in claim 22 wherein said applying said
2 final set of possible rules comprises selecting a highest priority rule in said final set of
3 possible rules.

1 24. The method as claimed in claim 20 wherein said range identifier
2 comprises an index value.

1 25. The method as claimed in claim 20 wherein said index values are
2 used to index into a N dimensional look-up table for a final rule.

1 26. The method as claimed in claim 19 wherein one of said N search
2 structures comprises a look-up table.

1 27. The method as claimed in claim 19 wherein one of said search
2 structures comprises a tree search structure.

ABSTRACT OF THE DISCLOSURE

As Internet packet flow increases, the demand for high speed packet filtering has grown. The present invention introduces a high-speed rule processing method that may be used for packet filtering. The method pre-processes a set of packet filtering rules such that the rules may be searched in parallel by a set of independent search units. Specifically, the rules are divided into N orthogonal dimensions that comprise aspects of each packet that may be examined and tested. Each of the N dimensions are then divided into a set of dimension rule ranges. Each rule range is assigned a value that specifies the rules that may apply in that range. The rule preprocessing is completed by creating a search structure to be used for classifying a packet into one of the rule ranges in each of the N dimensions. Each search structure may be used by an independent search unit such that all N dimensions may be searched concurrently. The packet processing method of the present invention activates the N independent search units to search the N pre-processor created search structures. The output of each of the N search structures is then logically combined to select a rule to be applied.

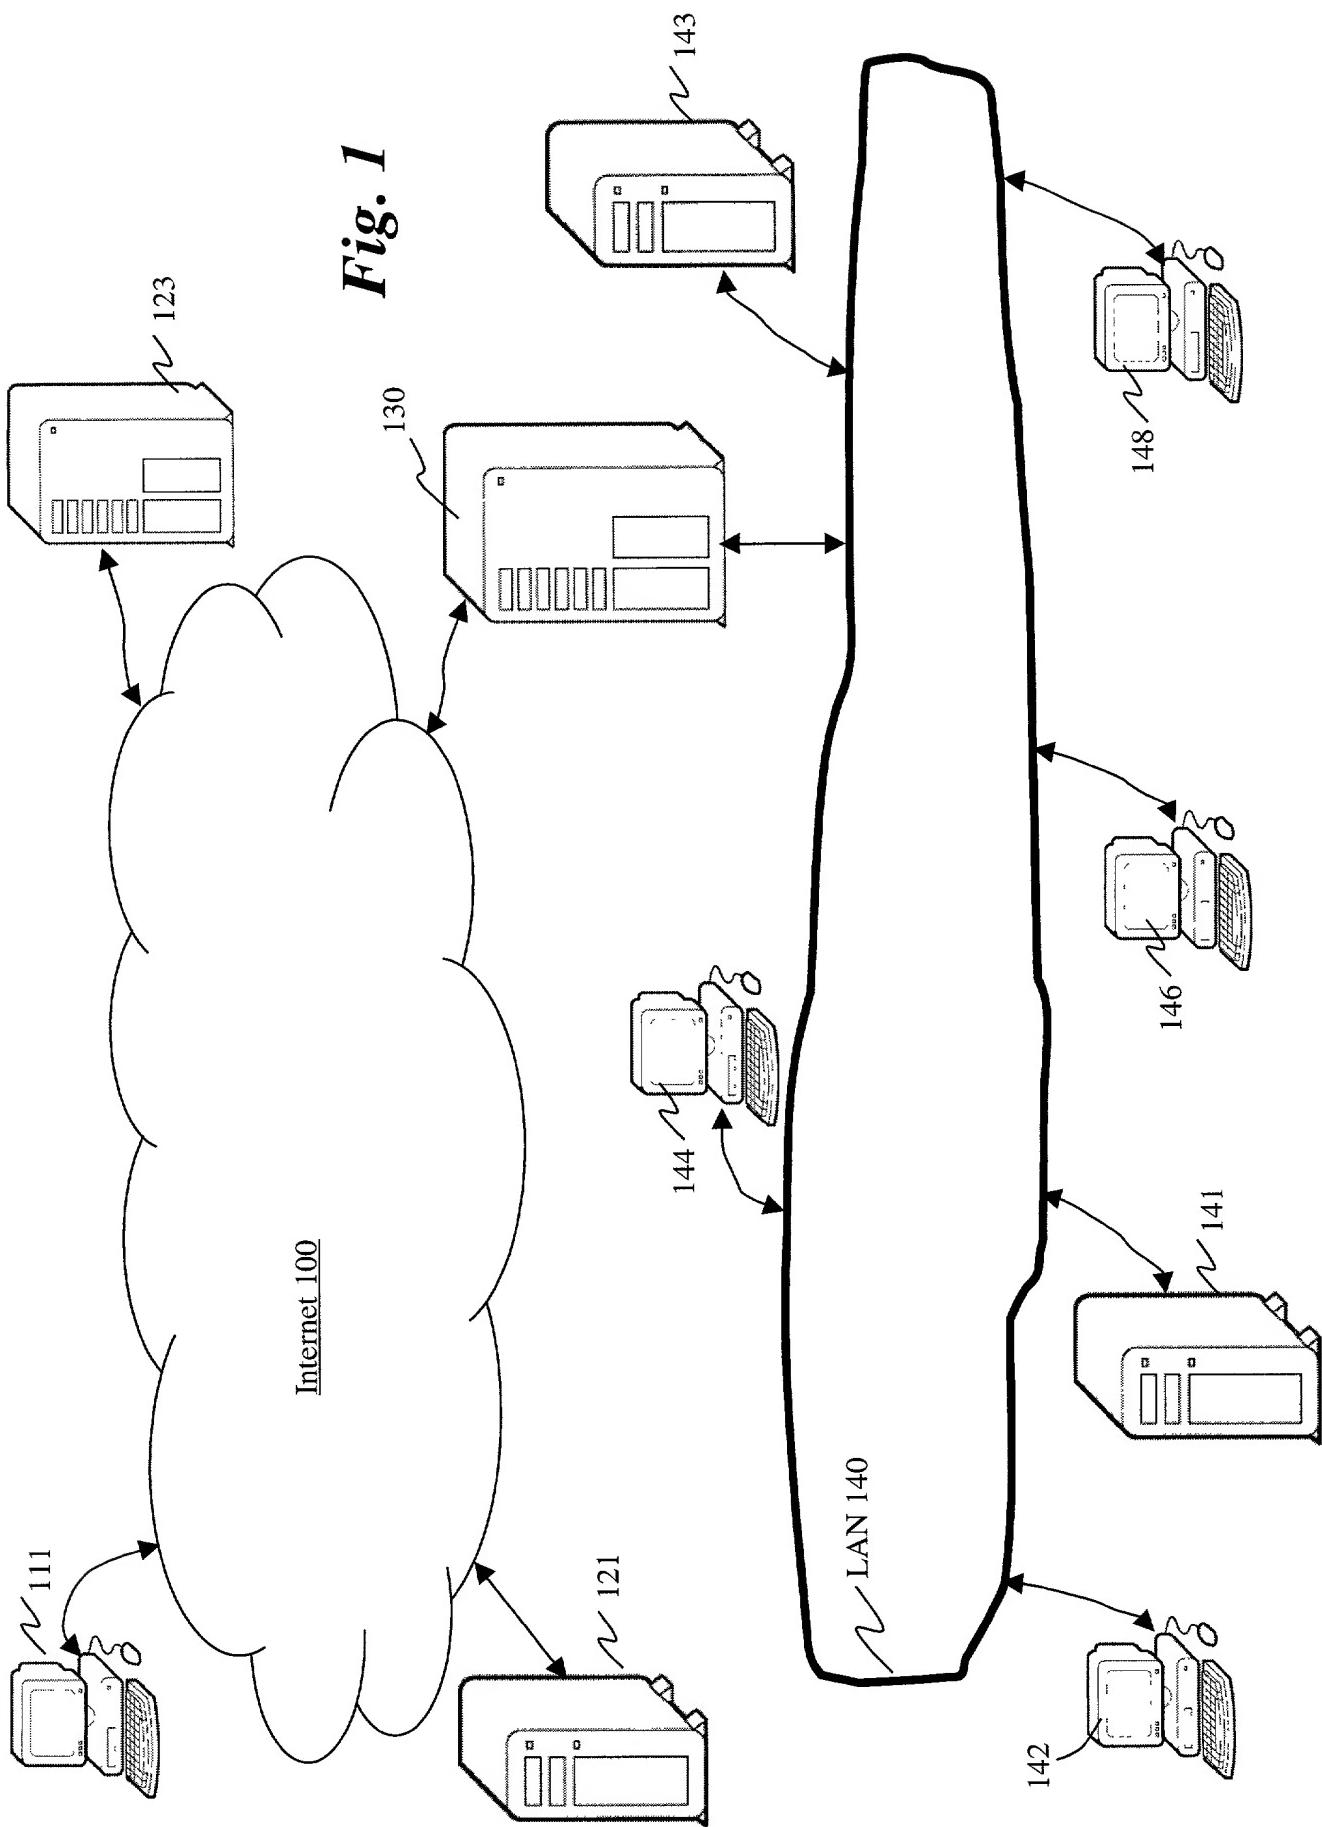


Fig. 2

Transport Control Protocol Packet

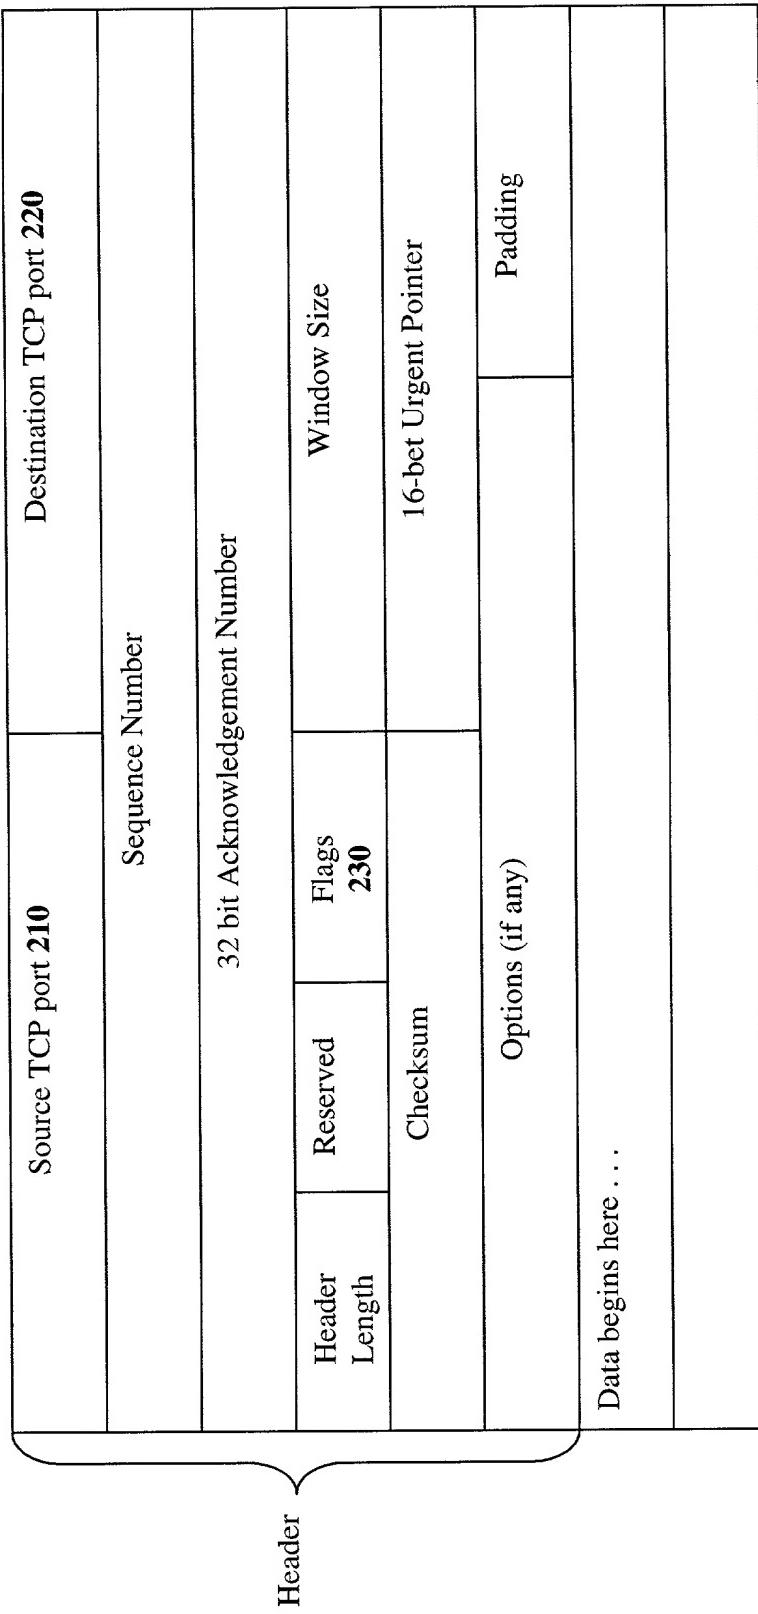


Fig. 3

Internet Protocol Packet

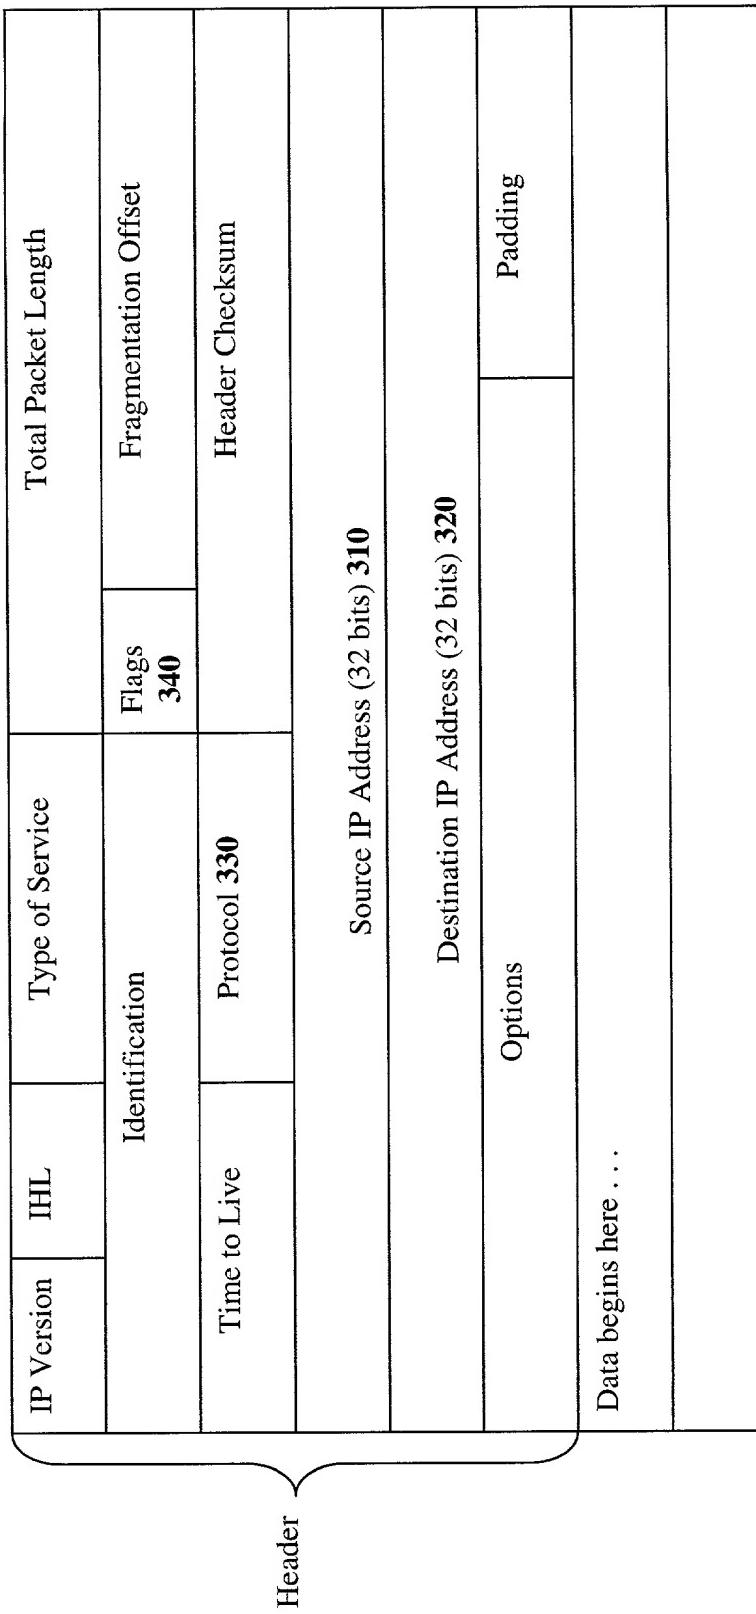


Fig. 4

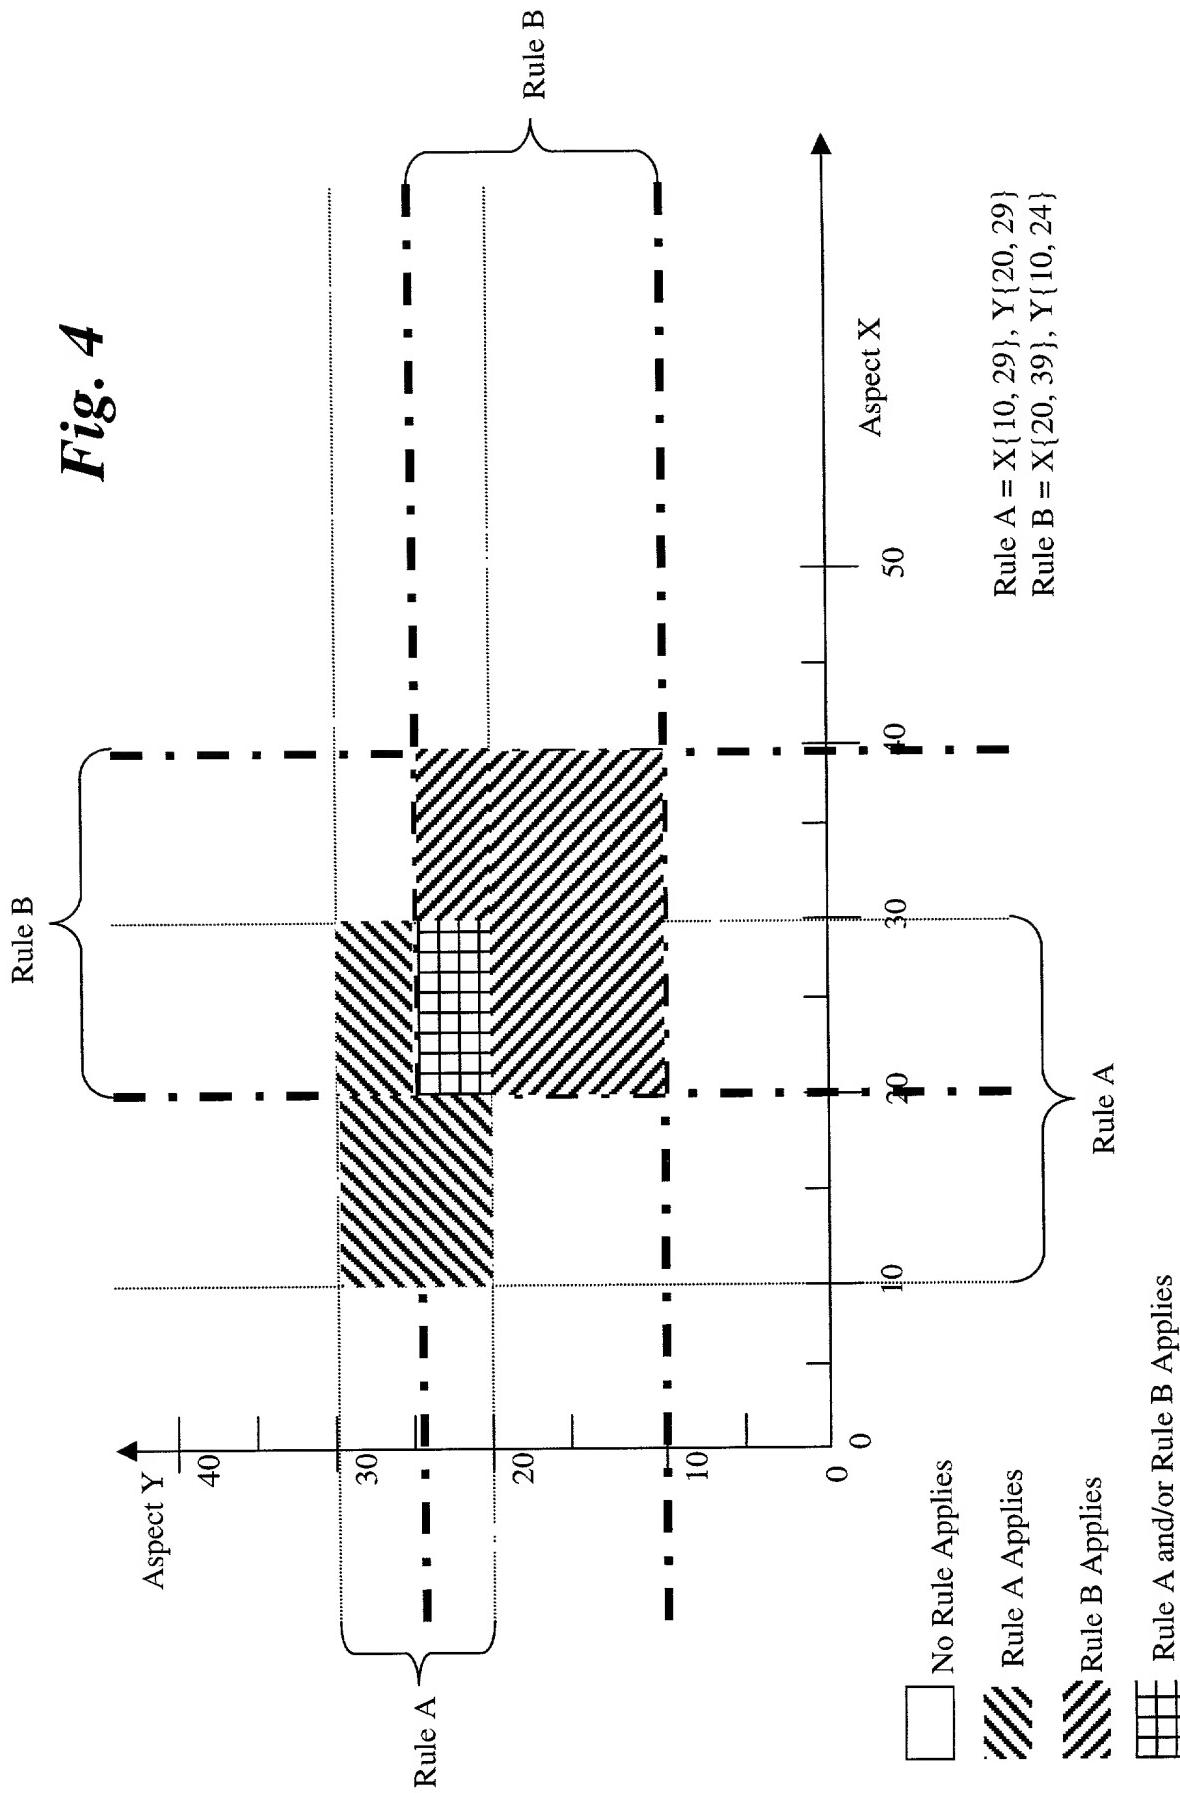


Fig. 5

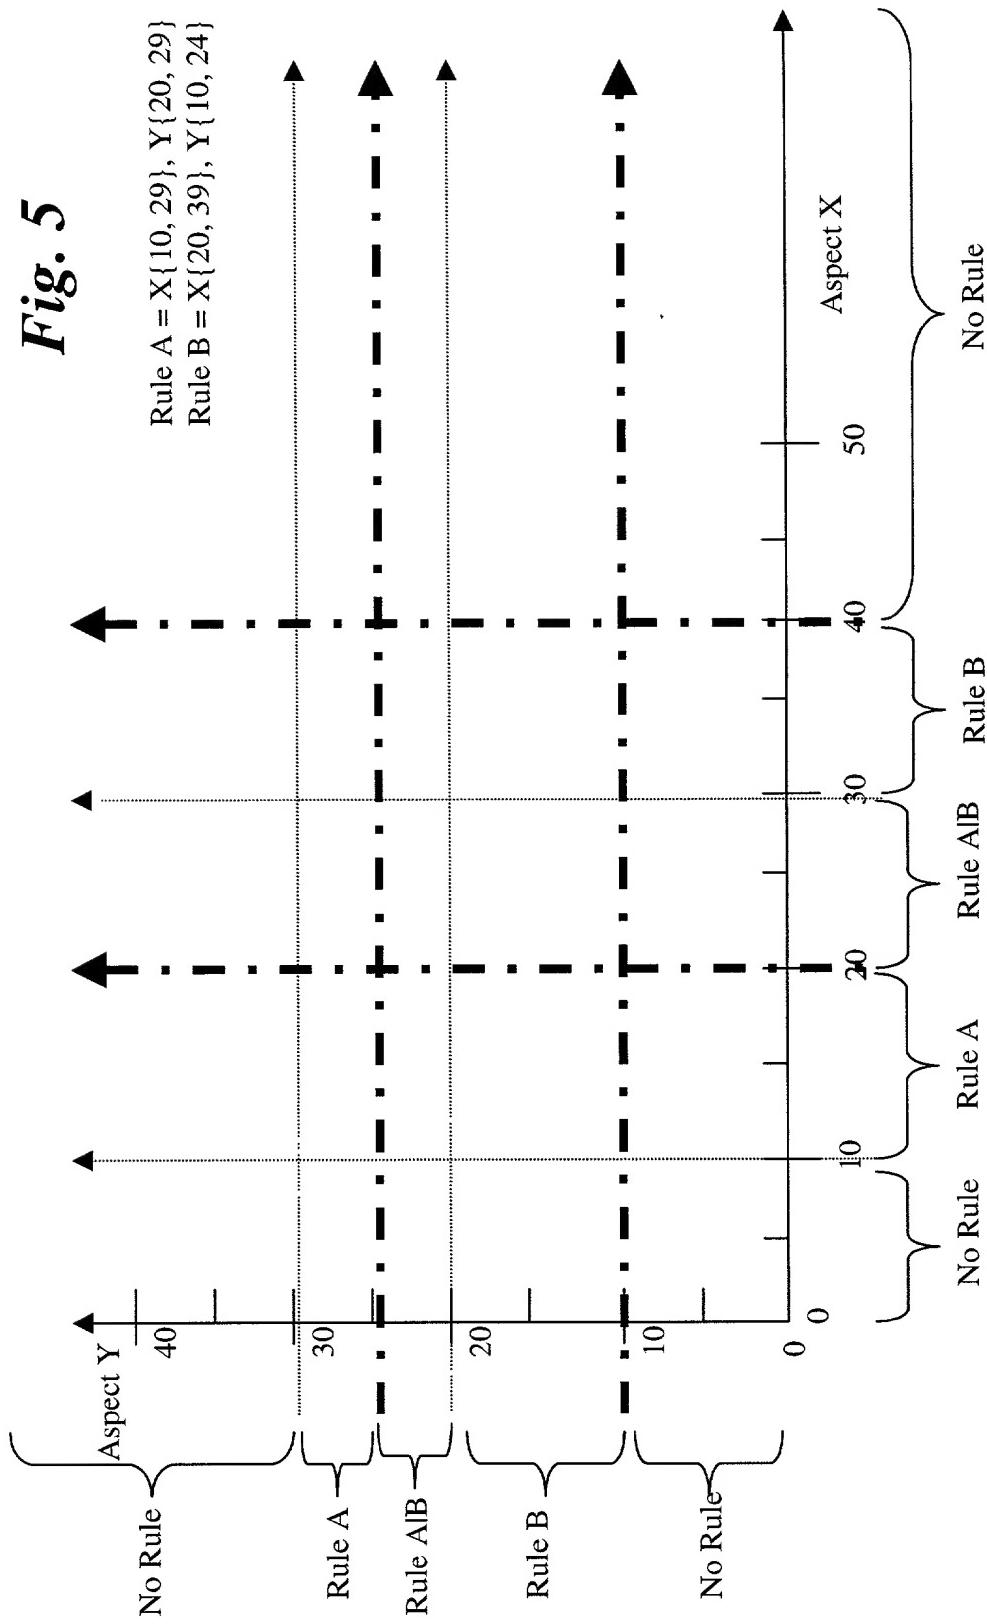


Fig. 6

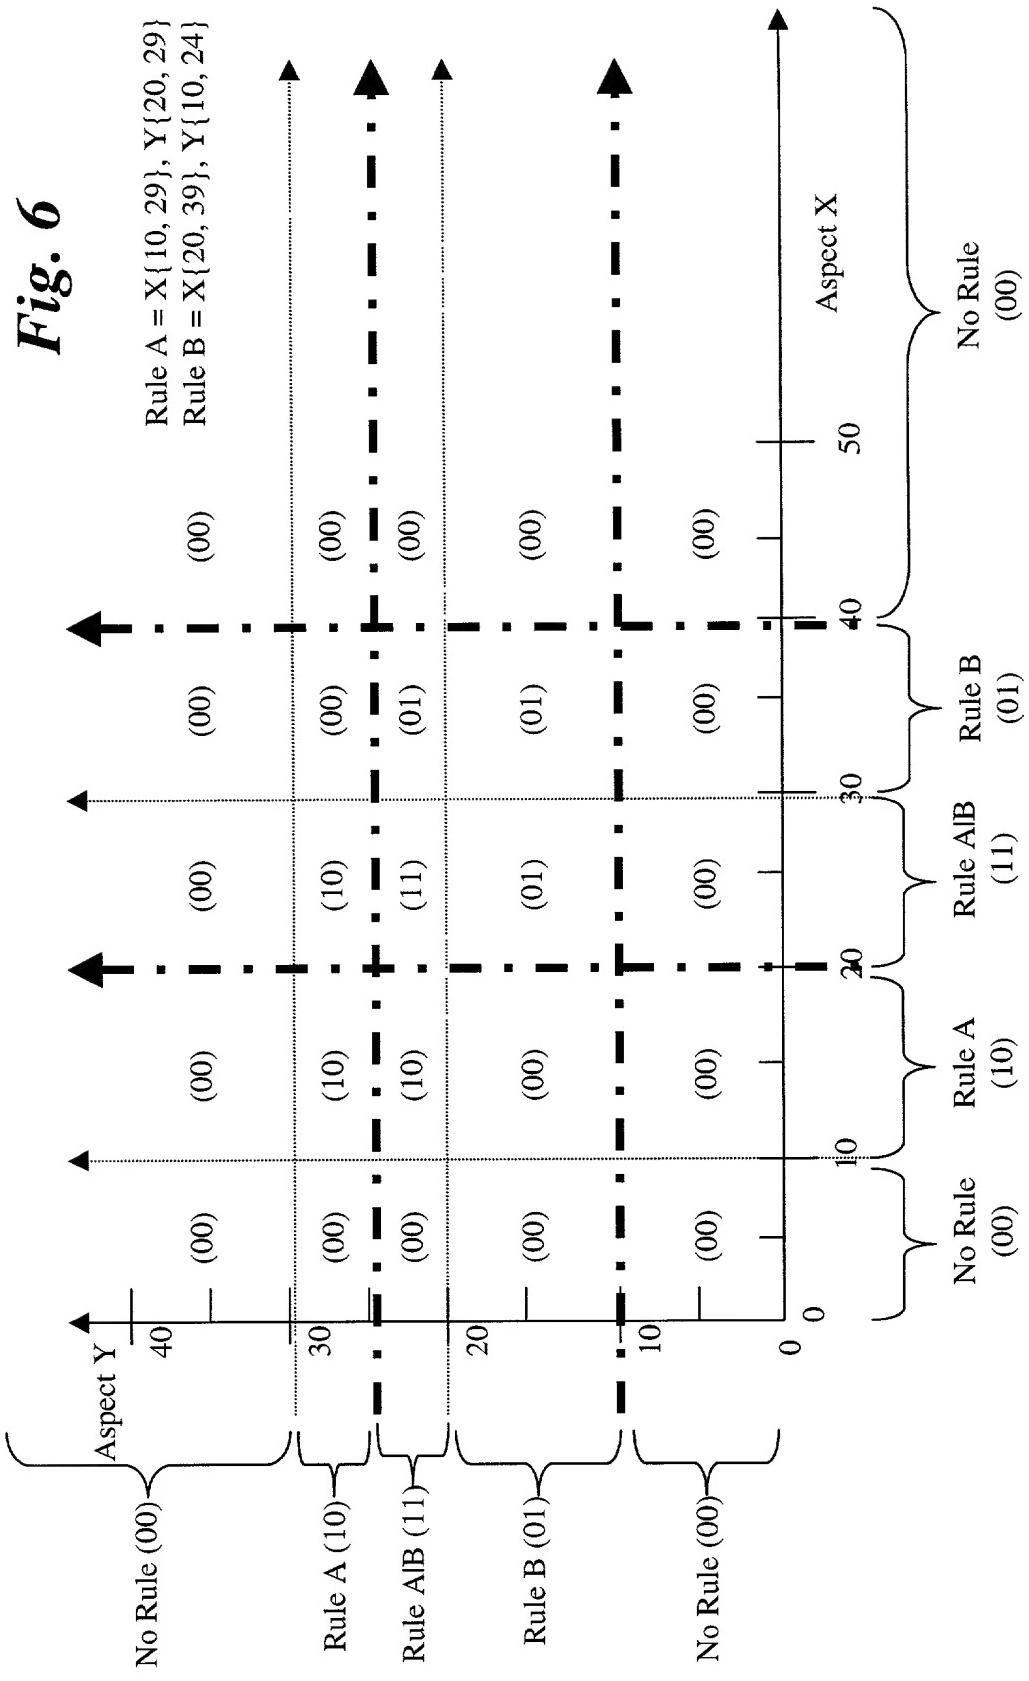


Fig. 7

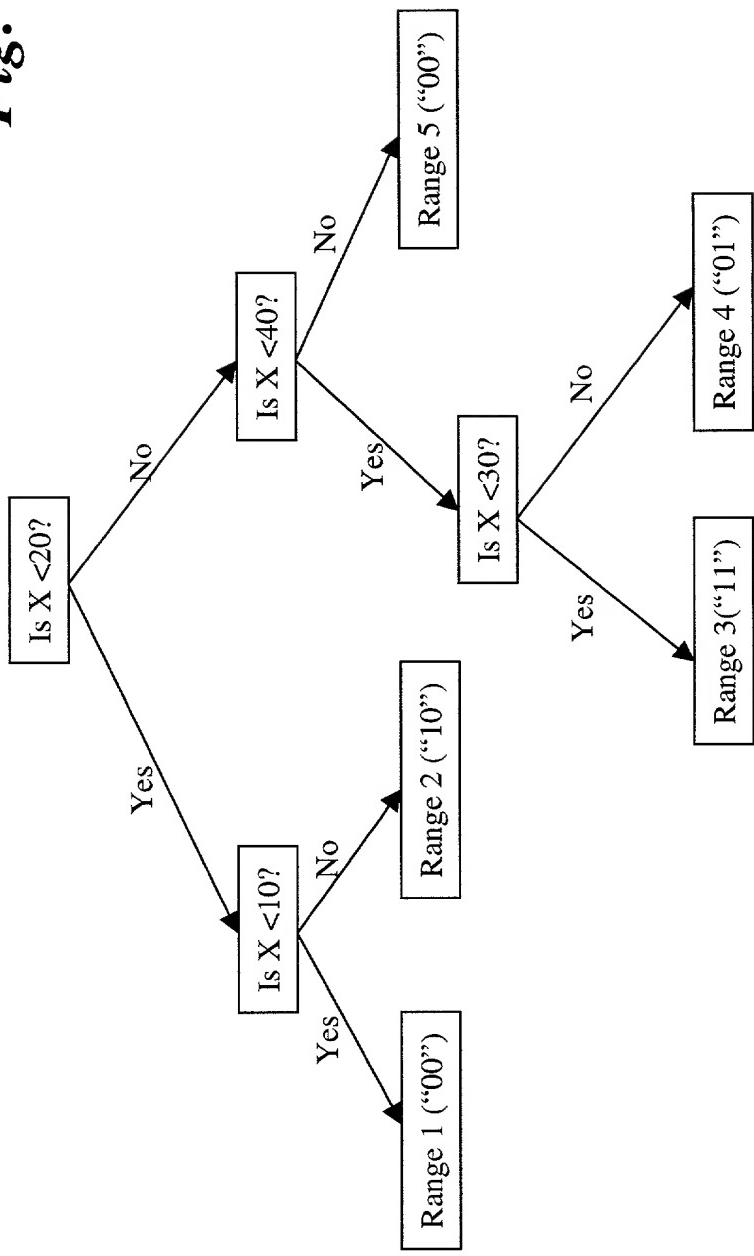
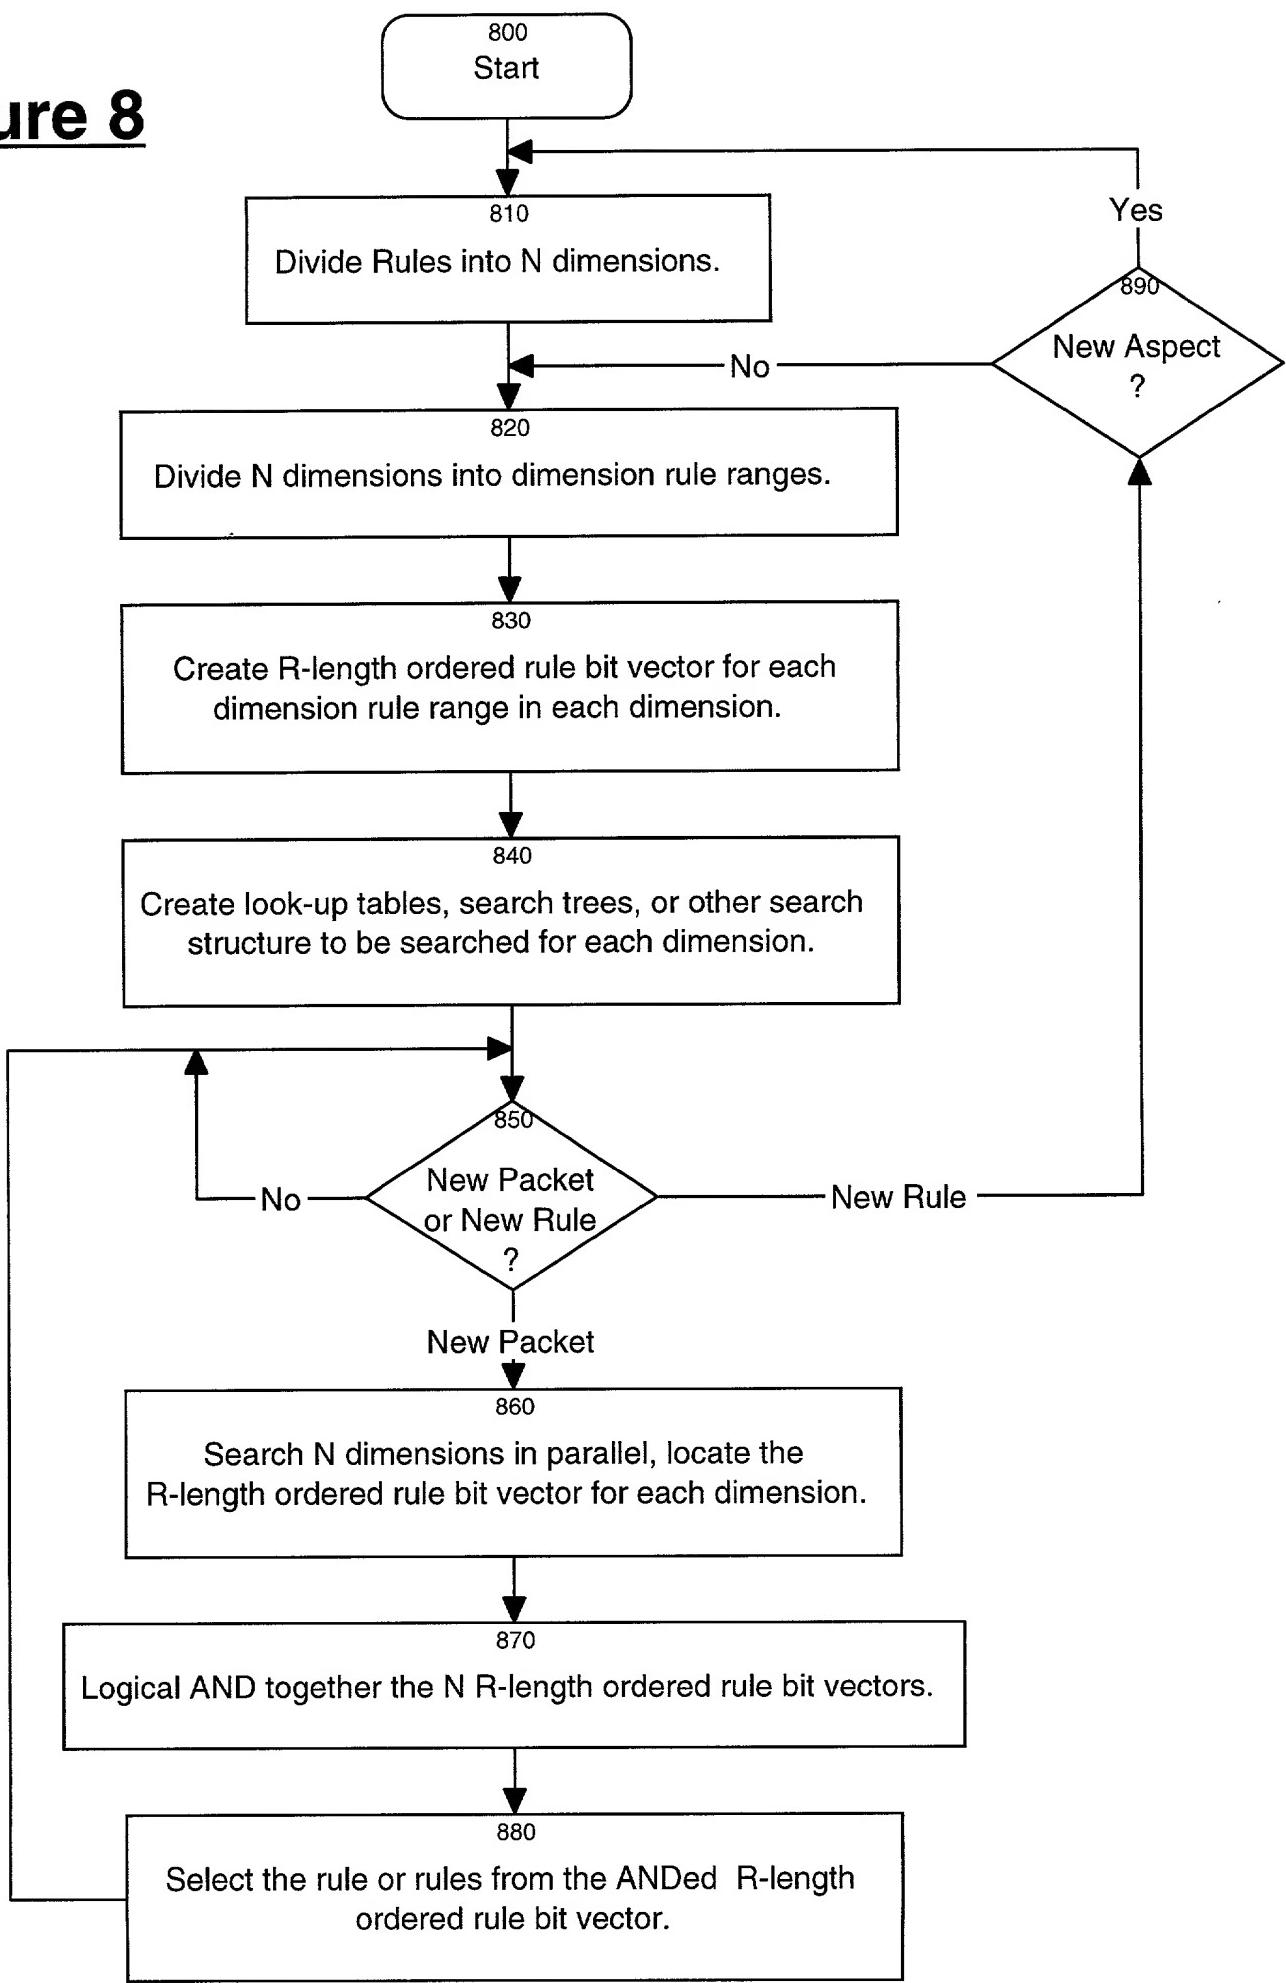


Figure 8



Attorney's Docket No.: 003813.P003

Patent

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

Method and Apparatus for High-Speed Network Rule Processing

the specification of which

XX is attached hereto.
____ was filed on _____ as

United States Application Number _____
or PCT International Application Number _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>	<u>Priority Claimed</u>			
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby claim the benefit under title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below

(Application Number)	Filing Date
(Application Number)	Filing Date

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
(Application Number)	Filing Date	(Status -- patented, pending, abandoned)

I hereby appoint Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; Amy M. Armstrong, Reg. No. P42,265; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. P41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Kent M. Chen, Reg. No. 39,630; Lawrence M. Cho, Reg. No. 39,942; Yong S. Choi, Reg. No. P43,324; Thomas M. Coester, Reg. No. 39,637; Roland B. Cortes, Reg. No. 39,152; Barbara Bokanov Courtney, Reg. No. P42,442; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Tarek N. Fahmi, Reg. No. 41,402; James Y. Go, Reg. No. 40,621; Richard Leon Gregory, Jr., P42,607; Dinu Gruia, Reg. No. P42,996; David R. Halvorson, Reg. No. 33,395; Thomas A. Hassing, Reg. No. 36,159; Phuong-Quan Hoang, P41,839; Willmore F. Holbrow III, Reg. No. P41,845; George W Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; Dag H. Johansen, Reg. No. 36,172; William W. Kidd, Reg. No. 31,772; Tim L. Kitchen, Reg. No. P41,900; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. P42,879; Darren J. Milliken, P42,004; Thinh V. Nguyen, P42,034; Kimberley G. Nobles, Reg. No. 38,255; Michael A. Proksch, Reg. No. P43,021; Babak Redjaian, P42,096; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Anand Sethuraman, Reg. No. P43,351; Charles E. Shemwell, Reg. No. 40,171; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Allan T. Sponseller, Reg. No. 38,318; Geoffrey T. Staniford, P43,151; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. P42,179; Edwin H. Taylor, Reg. No. 25,129; George G. C. Tseng, Reg. No. 41,355; Lester J. Vincent, Reg. No. 31,460; John Patrick Ward, Reg. No. 40,216; Stephen Warhola, Reg. No. P43,237; Ben J. Yorks, Reg. No. 33,609; and Norman Zafman, Reg. No. 26,250; my attorneys, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and James R. Thein, Reg. No. 31,710, my patent attorney; with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Full Name of Sole/First Inventor Raghunath Iyer

Inventor's Signature _____ Date _____

Residence Los Altos, CA (City, State) Citizenship India (Country)

Post Office Address 2260 Homestead Court, #309
Los Altos, CA 94024

Full Name of Second/Joint Inventor Sundar Iyer

Inventor's Signature _____ Date _____

Residence Palo Alto, CA _____ Citizenship India _____
(City, State) (Country)

Post Office Address 1103A Lyman House
Palo Alto, CA 94305

Full Name of Third/Joint Inventor Moti Jiandani

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Full Name of Fourth/Joint Inventor Ramana Rao

Inventor's Signature _____ Date _____

Residence _____ Citizenship _____
(City, State) (Country)

Post Office Address _____

Title 37, Code of Federal Regulations, Section 1.56
Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclosure information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclosure all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.